EXCEPTIONAL PARAMETERS FOR GENERIC A-HYPERGEOMETRIC SYSTEMS

LAURA FELICIA MATUSEVICH

ABSTRACT. The holonomic rank of an A-hypergeometric system $H_A(\beta)$ is conjectured to be independent of the parameter vector β if and only if the toric ideal I_A is Cohen Macaulay. We prove this conjecture in the case that I_A is generic by explicitly constructing more than $\mathrm{vol}\,(A)$ many linearly independent hypergeometric functions for parameters β coming from embedded primes of certain initial ideals of I_A.

1. INTRODUCTION

A-hypergeometric systems are systems of linear partial differential equations with polynomial coefficients that can be built out of a toric ideal and a parameter vector. Homogeneous toric ideals are themselves built out of combinatorial data: n distinct integer points lying in a hyperplane off the origin in d-dimensional space. We may assume that these points are the columns of a $d \times n$ integer matrix whose first row is made up of ones.

Definition 1.1. A $d \times n$ matrix A whose columns are distinct elements of $\{1\} \times \mathbb{Z}^{d-1}$ and generate $\mathbb{Z}^d$ as a lattice is said to be **homogeneous**. We set $m = n - d$.

As we have already mentioned, we will think of A as a point configuration in d-space. These points (the columns of A) will be called $a_1, \ldots, a_n$. The convex hull of $\{a_1, \ldots, a_n\}$, $\mathrm{conv}(A)$, is a $(d-1)$-dimensional polytope, whose normalized volume we denote by $\mathrm{vol}\,(A)$.

Definition 1.2. Given a homogeneous matrix A, the **toric ideal** I_A is the ideal of the polynomial ring $\mathbb{C}[\partial_1, \ldots, \partial_n]$ given by

$$I_A = \langle \partial^u - \partial^v : u, v \in \mathbb{N}^n, A \cdot u = A \cdot v \rangle \ .$$

Here we use multi-index notation $\partial^u = \partial_1^{u_1} \cdots \partial_n^{u_n}$.

A-hypergeometric systems are non-commutative objects, they are left ideals in the Weyl algebra D. The Weyl algebra is the quotient of the free associative algebra with generators $x_1, \ldots, x_n, \partial_1, \ldots, \partial_n$, modulo the relations

$$x_i x_j = x_j x_j \ , \ \partial_i \partial_j = \partial_j \partial_i \ , \ \partial_i x_j = x_j \partial_i + \delta_{ij} \ , \ 1 \leq i, j \leq n \ ,$$

where δ_{ij} is the Kronecker delta.

The author was partially supported by a Julia B. Robinson fellowship at UC Berkeley.

Definition 1.3. Given a homogeneous matrix A and a vector $\beta \in \mathbb{C}^d$, the **A-hypergeometric system with parameter vector** β is the left ideal in the Weyl algebra generated by I_A and the homogeneity operators:

$$(\sum_{j=1}^{n} a_{ij} x_j \partial_j) - \beta_i \ , \ 1 \leq i \leq d \ .$$

Hypergeometric systems were introduced in the late eighties by Gel′fand, Kapranov and Zelevinsky (see, for instance, [4]). As we define them here, they are regular holonomic. This means two things. First, that the *holonomic rank* of $H_A(\beta)$, that is, the dimension of the space of holomorphic solutions of $H_A(\beta)$ around a nonsingular point, is well defined, and second, that the solutions of $H_A(\beta)$ can be represented as power series with logarithms. The holonomic rank of $H_A(\beta)$ is denoted $\operatorname{rank}(H_A(\beta))$. One of the first results shown by Gel′fand, Kapranov and Zelevinsky about A-hypergeometric systems is that, when the underlying toric ideal I_A is Cohen Macaulay, $\operatorname{rank}(H_A(\beta)) = \operatorname{vol}(A)$ for all $\beta \in \mathbb{C}^d$. A proof of this theorem can be found in [13, Section 4.3]. This equality can fail if I_A is not Cohen Macaulay.

However, as is shown in [1], [13, Theorem 3.5.1,Equation 4.3], if we drop the Cohen Macaulayness hypothesis, $\operatorname{rank}(H_A(\beta)) \geq \operatorname{vol}(A)$ for all $\beta \in \mathbb{C}^d$, and equality holds for generic β.

We define the *exceptional set* of A to be:

$$\mathcal{E}(A) := \{\beta \in \mathbb{C}^d : \operatorname{rank}(H_A(\beta)) > \operatorname{vol}(A)\} \ .$$

Elements of $\mathcal{E}(A)$ are called *exceptional parameters*. The following conjecture, due to Bernd Sturmfels, relates the existence of exceptional parameters and the Cohen Macaulayness of I_A.

Conjecture 1.4. *The exceptional set of A is empty if and only if I_A is Cohen Macaulay.*

The "if" part of this conjecture is the aforementioned result by Gel′fand, Kapranov and Zelevinsky. As for the "only if" part, it has been proved in various special cases. The $d = 2$ case of Conjecture 1.4 is a result of Cattani, D'Andrea and Dickenstein (see [3]). The $n - d = 2$ case was proved by the author in [7]. Finally, Saito has shown that Conjecture 1.4 is true when the convex hull of the configuration A is a simplex (see [12]).

In this article, we study Conjecture 1.4 in the case that I_A is a generic toric ideal.

Definition 1.5. A lattice ideal is **generic** if it has a minimal generating set of binomials with full support.

Generic lattice ideals were introduced by Irena Peeva and Bernd Sturmfels in [10]. One of the results in that article is that "most" toric ideals are generic.

The following is the main result in this article.

Theorem 1.6. *Let I_A be a generic non Cohen Macaulay toric ideal. Then the exceptional set of A, $\mathcal{E}(A)$ contains an affine space of dimension $d - 2$. In particular, $\mathcal{E}(A) \neq \emptyset$.*

In order to prove this theorem, we will need a characterization of Cohen Macaulayness for generic toric ideals. We will use Lemma 4.1, that relates this notion to the existence of embedded primes of certain initial ideals of our toric ideal.

This article is organized as follows. Sections 2 and 3 contain background material on toric algebra and hypergeometric functions respectively. In Sections 4, 5, 6, and 7, we prove Theorem 1.6. Section 8 contains a completely worked out example, and in Section 9 we show that, under no hypotheses on A, the exceptional set $\mathcal{E}(A)$ is Zariski constructible.

2. Standard pairs of initial ideals of toric ideals

In this section we introduce the notion of standard pairs. These objects play a fundamental role in the study of the associated primes of monomial ideals: if (∂^η, σ) is a standard pair of a monomial ideal $M \subset \mathbb{C}[\partial_1, \ldots, \partial_n]$, then $\langle \partial_j : j \notin \sigma \rangle$ is an associated prime of M. Moreover, all associated primes of M arise this way.

In the special case of initial ideals of toric ideals, standard pairs admit a polyhedral description (Theorem 2.2).

Definition 2.1. Let M be a monomial ideal of $\mathbb{C}[\partial_1, \ldots, \partial_n]$. A **standard pair** of M is a pair (∂^η, σ), where $\eta \in \mathbb{N}^n$ and $\sigma \subset \{1, \ldots, n\}$ subject to the following three conditions:

1. $\eta_i = 0$ for $i \in \sigma$;
2. For all choices of integers $\mu_i \geq 0$, $i \in \sigma$, the monomial $\partial^\eta \cdot \prod_{i \in \sigma} \partial_i^{\mu_i}$ is not in M.
3. For all $l \notin \sigma$, there exist $\mu_i \geq 0$, $i \in \sigma \cup \{l\}$, such that $\partial^\eta \cdot \partial_l^{\mu_l} \cdot \prod_{i \in \sigma} \partial_i^{\mu_i}$ is in M.

The set of standard pairs of M is denoted $S(M)$. A standard pair (∂^η, σ) such that the ideal $\langle \partial_i : i \notin \sigma \rangle$ is a minimal associated prime of M is called **top-dimensional**. Standard pairs that are not top-dimensional are called **embedded**.

If A is a homogeneous $d \times n$ matrix and $w \in \mathbb{R}^n$ is a generic weight vector for I_A, that is, $\text{in}_w(I_A)$ is a monomial ideal, we can study $S(\text{in}_w(I_A))$ using combinatorial techniques. For instance, a standard pair of $\text{in}_w(I_A)$ is top-dimensional if and only if the cardinality of $\{1, \ldots, n\} \backslash \sigma$ is equal to m (see [5, Corollary 2.9]). The following polyhedral characterization of standard pairs of initial ideals of toric ideals is due to Serkan Hoşten and Rekha Thomas (see [5, Theorems 2.3, 2.5]). Choose a $\mathbb{Z}$-basis of $\ker_{\mathbb{Z}}(A)$ and form an $n \times m$ matrix $B = (b_{ij})$ whose columns are the vectors in this basis. This matrix B is called a *Gale dual* of A.

Theorem 2.2. *A pair (∂^η, σ), where $\eta \in \mathbb{N}^n$ and $\eta_i = 0$ for $i \in \sigma$, is a standard pair of the monomial ideal $\text{in}_w(I_A)$ if and only if 0 is the only lattice point in the polytope*

$$P_\eta^{\bar{\sigma}} := \{ y \in \mathbb{R}^m : (B \cdot y)_j \leq \eta_j \ , j \notin \sigma \ ; -(w)^t (B \cdot y) \leq 0 \} \ ,$$

and all the inequalities $(B \cdot y)_j \leq \eta_j \ , j \notin \sigma$, are essential, that is, removing an inequality introduces a new lattice point z into the resulting polyhedron. We may assume that z is such that $-w^t(B \cdot z)$ is strictly negative.

The fact that $P_\eta^{\bar{\sigma}}$ is a polytope (and hence a bounded set) has the following linear algebra consequence.

Lemma 2.3. *Let (∂^η, σ) be a standard pair of $\operatorname{in}_w(I_A)$, where w is such that this is a monomial ideal. Then the set $\{(b_{i1}, \dots, b_{im}) : i \notin \sigma\}$ contains a linearly independent subset of cardinality m.*

Proof. By contradiction, suppose that no subset of cardinality m of $\{(b_{i1}, \dots, b_{im}) : i \notin \sigma\}$ is linearly independent. This means that the matrix whose rows are the rows of B indexed by $i \notin \sigma$ has rank strictly less than m. Consequently, we can find rational numbers $s_1, \dots, s_m$ not all zero, such that $(B \cdot (s_1, \dots, s_m)^t)_i = 0$, for all $i \notin \sigma$. But then at least half of the line $\{\lambda(s_1, \dots, s_m)^t : \lambda \in \mathbb{R}\}$ is contained in $P_\eta^{\bar{\sigma}}$, contradicting that this set is bounded. □

Associated primes of initial ideals of toric ideals come in saturated chains. This is the content of the following theorem, due to Hoşten and Thomas [5, Theorem 3.1].

Theorem 2.4. *Let I_A be a toric ideal and $w \in \mathbb{R}^n$ a generic weight vector for I_A. If $\mathfrak{p}$ is an embedded prime of $\operatorname{in}_w(I_A)$, then $\mathfrak{p}$ contains an associated prime $\mathfrak{q}$ of $\operatorname{in}_w(I_A)$ such that $dim(\mathfrak{q}) = dim(\mathfrak{p}) + 1$.*

Finally, we include here a combinatorial lemma that will be useful later on. I am grateful to Bernd Sturmfels, who provided this beautiful proof.

Lemma 2.5. *Let $c^1, \dots, c^{m+2} \in \mathbb{R}^m$ and $k_1, \dots k_{m+2} \in \mathbb{R}$ and consider $P = \{z \in \mathbb{R}^m : c^j \cdot z \leq k_j, j = 1, \dots m+2\}$. Suppose that this set is nonempty. There is a set $T \subset \{1, \dots, m+1\}$ of cardinality m such that, for each $j \in T$, the set obtained from P by reversing the inequality $c^j \cdot z \leq k_j$ is unbounded. Moreover, the set $\{c^j : j \in T\}$ is linearly independent.*

Proof. We use the technique of Gale duality, as introduced in [16, Chapter 6]. Consider the vector configuration $\mathcal{C} = \{c_1, \dots, c_{m+2}\} \subset \mathbb{R}^m$. Let $\mathcal{B} = \{b_1, \dots, b_{m+2}\} \subset \mathbb{R}^2$ be its Gale dual configuration. We define $\mathcal{N}$ as the set of indices j such that reversing the j-th inequality in P produces a bounded set. Now, $i \in \mathcal{N}$ if an only if the configuration $\mathcal{C} \backslash \{c_i\}$ is totally cyclic, that is, if and only if there exists a positive linear dependence among the elements of $\mathcal{C} \backslash \{c_i\}$.

Using results from [16, Sections 6.3(b), 6.4], namely that the dual operation to deletion is contraction, and that dual of a totally cyclic vector configuration is an acyclic vector configuration, we see that $i \in \mathcal{N}$ if and only if $\mathcal{B}/b_i \subset \mathbb{R}$ is acyclic, that is, if and only if there is a linear functional $c \in \mathbb{R}$ such that $c \cdot b > 0$ for all $b \in \mathcal{B}/b_i$. But this can only happen if b_i is an extreme ray of $\mathcal{B}$, which means that there exists $c' \in \mathbb{R}^2$ such that $c' \cdot b_i = 0$ and $c' \cdot b_j > 0$ for all $1 \leq j \leq m+2$, $j \neq i$. Since a configuration in $\mathbb{R}^2$ has at most two extreme rays, we conclude that the cardinality of $\mathcal{N}$ is at most 2. The last assertion of this lemma follows from the same arguments used to prove Lemma 2.3. □

3. Canonical Hypergeometric Series

In this section we review material about canonical logarithm-free hypergeometric series. Our source is [13, Sections 2.5, 3.1, 3.4, 4.1]. We start by relating standard pairs to hypergeometric functions through the concept of fake exponents.

Definition 3.1. Let A be a homogeneous $d \times n$ matrix, $\beta \in \mathbb{C}^d$, and $w \in \mathbb{R}^n$ a generic weight vector for I_A. The set of **fake exponents of $H_A(\beta)$ with respect to w** is the vanishing set of the following zero dimensional ideal of the (commutative) polynomial ring $\mathbb{C}[\theta_1, \ldots, \theta_n]$:

$$\bigcap_{(\partial^\eta, \sigma) \in S(\mathrm{in}_w(I_A))} \left(\langle \theta_i - \eta_i : i \notin \sigma \rangle + \langle A \cdot \theta - \beta \rangle \right),$$

where θ is the vector $(\theta_1, \ldots, \theta_n)^t$.

If A is homogeneous and $\beta \in \mathbb{C}^d$, the hypergeometric system $H_A(\beta)$ is regular holonomic. This means that the solutions of $H_A(\beta)$ can be written as power series with logarithms. Moreover, if w is a generic weight vector for I_A, we can use the techniques in [13, Sections 2.5, 3.4, 4.1] to build a basis of the solution space of $H_A(\beta)$, whose elements are called *canonical series* (with respect to w). This basis is defined by Proposition 3.2.

Notice that we can extend the notion of term order to the ring of power series with logarithms as follows:

$$x^\alpha \log(x)^\gamma \leq x^{\alpha'} \log(x)^{\gamma'} \iff \mathrm{Re}\,(w \cdot \alpha) \leq \mathrm{Re}\,(w \cdot \alpha').$$

Here we mean $x^\alpha = x^{\alpha_1} \cdots x^{\alpha_n}$ for $\alpha \in \mathbb{C}^n$, where $x_i^{\alpha_i} = \exp(\alpha_i \log(x_i))$; and $\log(x)^\gamma = \log(x_1)^{\gamma_1} \cdots \log(x_n)^{\gamma_n}$, for $\gamma \in \mathbb{N}^n$. If we refine this ordering lexicographically, we can define the initial term of a power series with logarithms φ (if it exists) as $\mathrm{in}_{\prec_w}(\varphi) =$ the *least* term of φ with respect to this ordering. By [13, Proposition 2.5.2], $\mathrm{in}_{\prec_w}(\varphi)$ exists when φ is a solution of $H_A(\beta)$ that converges in a certain region of $\mathbb{C}^n$ that depends on w.

Proposition 3.2. *There exist a basis of the solution space of $H_A(\beta)$ such that*

- *if φ is an element of this basis, there exists a fake exponent v of $H_A(\beta)$ with respect to w such that $\mathrm{in}_{\prec_w}(\varphi) = x^v \log(x)^\gamma$ for some $\gamma \in \mathbb{N}^n$,*
- *if φ and ψ are elements of this basis, then $\mathrm{in}_{\prec_w}(\varphi)$ is not a term appearing in ψ.*

This basis is called the **basis of canonical series with respect to** w.

Proof. This follows from Corollary 2.5.11, Corollary 3.1.6 and Lemma 4.1.3 in [13]. □

Remark 3.3. There might be fake exponents of $H_A(\beta)$ with respect to w that do not give initial monomials of canonical series (see Example 3.1.8 in [13]). Also the same fake exponent might give rise to two different canonical series. If this is the case, then at least one of those canonical series will have logarithms.

Using the results from [13, Section 3.4] we can pinpoint exactly which fake exponents give initial monomials of *logarithm-free* canonical series solutions of $H_A(\beta)$. To do this we need to introduce the following concepts. The *negative support* of a vector $v \in \mathbb{C}^n$ is the set:

$$\mathrm{nsupp}(v) := \{i \in \{1, \ldots, n\} : v_i \in \mathbb{Z}_{<0}\}.$$

A vector $v \in \mathbb{C}^n$ has *minimal negative support* with respect to A if

$$u \in \ker_{\mathbb{Z}}(A) \text{ and } \mathrm{nsupp}(v-u) \subseteq \mathrm{nsupp}(v) \text{ imply}$$
$$\mathrm{nsupp}(v-u) = \mathrm{nsupp}(v)\,.$$

In this case, let

$$N_v = \{u \in \ker_{\mathbb{Z}}(A) : \mathrm{nsupp}(v-u) = \mathrm{nsupp}(v)\}\,,$$

and define the following formal power series:

$$\phi_v = \sum_{-u \in N_v} \frac{[v]_{u_-}}{[u+v]_{u_+}} x^{v+u} \quad , \tag{1}$$

where

$$[v]_{u_-} = \prod_{i:u_i<0} \prod_{j=1}^{-u_i} (v_i - j + 1) \quad \text{and} \quad [u+v]_{u_+} = \prod_{i:u_i>0} \prod_{j=1}^{u_i} (v_i + j)\,.$$

Theorem 3.4. [13, Theorem 3.4.14, Corollary 3.4.15] *Let $v \in \mathbb{C}^n$ be a fake exponent of $H_A(\beta)$ with minimal negative support. Then the series ϕ_v defined in (1) is a canonical solution of the A-hypergeometric system $H_A(\beta)$. In particular, ϕ_v converges in a region of $\mathbb{C}^n$. The set:*

$$\{\phi_v : v \textit{ is a fake exponent with minimal negative support}\}$$

is a basis of the space of logarithm-free solutions of $H_A(\beta)$.

We now mention a way to distinguish which vectors with minimal negative support are fake exponents.

Proposition 3.5. *Let $v \in \mathbb{C}^n$ with minimal negative support such that $A \cdot v = \beta$. Then v is a (fake) exponent of $H_A(\beta)$ with respect to a weight vector w if and only if*

$$w \cdot v = \mathit{min}\{w \cdot u : \mathrm{nsupp}(u) = \mathrm{nsupp}(v) \textit{ and } u - v \in \ker_{\mathbb{Z}}(A)\}$$

and this minimum is attained uniquely.

At the moment, a detailed characterization such as we have for logarithm-free A-hypergeometric series does not exist for logarithmic A-hypergeometric functions. However, information about logarithmic series will be necessary to prove Theorem 1.6. We reproduce two results about logarithmic hypergeometric functions. The first is an observation from [7].

Observation 3.6. Let ψ be a solution of $H_A(A \cdot v)$. This function is of the form:

$$\psi = \sum c_{\alpha,\gamma} x^\alpha \log(x)^\gamma\,,$$

where the sum runs over α such that $A\alpha = A{\cdot}v$, $\gamma \in \{0, 1, \ldots, h-1\}^n$ for $h = \mathrm{rank}\,(H_A(A{\cdot}v))$.

The set $\mathcal{S} := \{\gamma \in [0, h-1]^n \cap \mathbb{N}^n : \exists \alpha \in \mathbb{C}^n \text{ such that } c_{\alpha,\gamma} \neq 0\}$ is partially ordered with respect to:

$$(\gamma_1, \ldots, \gamma_n) \leq (\gamma'_1, \ldots, \gamma'_n) \Longleftrightarrow \gamma_i \leq \gamma'_i,\ i = 1, \ldots, n\,.$$

Denote by $\mathcal{S}_{\max}$ the set of maximal elements of $\mathcal{S}$. Let $\delta \in \mathcal{S}_{\max}$ and $f_\delta = \sum_{\alpha \in \mathbb{C}^n} c_{\alpha,\delta} x^\alpha$. Write

$$\psi = \psi_\delta + \log(x)^\delta f_\delta ,$$

so that the logarithmic terms in ψ_δ are either less than or incomparable to δ. If P is a differential operator that annihilates ψ, we have:

$$0 = P\psi = P\psi_\delta + \log(x)^\delta P f_\delta + \text{terms whose log factor is lower than } \delta.$$

Since $P\psi_\delta$ is a sum of terms whose log factor is either lower than δ or incomparable to δ, we conclude that Pf_δ must be zero. This implies that f_δ is a logarithm-free A-hypergeometric function of degree $A \cdot v$. *Moreover, if $\partial_1 \psi$ is logarithm-free, then $\partial_1 f_\delta$ must vanish.*

The following result is due to Saito (see [12]).

Proposition 3.7. *Let $\varphi = \sum x^u g_u(\log(x))$ be a solution of $H_A(\beta)$. Then the polynomials g_u are of the form:*

$$g_u(t_1, \ldots, t_n) = c_0 + c^{(1,1)} \cdot (t_1, \ldots, t_n) + \cdots + \prod_{i=1}^{l} c^{(l,i)} \cdot (t_1, \ldots, t_n),$$

where $c_0 \in \mathbb{C}$, the vectors $c^{(i,j)}$ belong to the kernel of A, and

$$c^{(i,j)} \cdot (t_1, \ldots, t_n) = c_1^{(i_j)} t_1 + \cdots + c_n^{(i,j)} t_n \ .$$

Moreover, if $g_u \neq 0$, then $A \cdot u = \beta$.

Finally, the techniques from [13, Section 3.5] imply the following proposition.

Proposition 3.8. *Let $\beta, \beta' \in \mathbb{C}^n$ and suppose that* $\operatorname{rank}(H_A(\beta + \epsilon\beta')) \geq t$ *for* $0 < \epsilon < \epsilon_0$. *Then* $\operatorname{rank}(H_A(\beta)) \geq t$.

4. Constructing exceptional parameters

In this section we start working towards the proof of Theorem 1.6. To do this, we will use the following characterization of the Cohen Macaulay property for generic toric ideals via associated primes of reverse lexicographic initial ideals of I_A.

Lemma 4.1. *Let I_A be a generic toric ideal. Then* $\text{in}_{-e_i}(I_A)$ *is a monomial ideal for all $1 \leq i \leq n$. Moreover, for a generic toric ideal, the following are equivalent:*

1. *I_A is Cohen Macaulay,*
2. *For all i,* $\text{in}_{-e_i}(I_A)$ *is free of embedded primes,*
3. *For all i,* $\text{in}_{-e_i}(I_A)$ *is Cohen Macaulay.*

Proof. The first assertion follows from [11, Lemma 8.4]. To see that (1) implies (2), notice that, since I_A is a Cohen Macaulay prime ideal , $I_A + \langle \partial_i \rangle$ is Cohen Macaulay, and thus free of embedded primes, for all $1 \leq i \leq n$. As a consequence, for each i, the ideal $\text{in}_{-e_i}(I_A)$, which is the ideal of $\mathbb{C}[\partial_1, \ldots \partial_n]$ generated by $(I_A + \langle \partial_i \rangle) \cap \mathbb{C}[\partial_1, \ldots, \hat{\partial}_i, \ldots, \partial_n]$, is free of embedded primes. The second implication follows from Theorems 2.5 and 3.1 in [9]. Finally, $\text{in}_{-e_i}(I_A)$ being Cohen Macaulay implies that I_A is Cohen Macaulay. $\square$

Remark 4.2. From now on, we assume that I_A is a generic non Cohen Macaulay toric ideal.

Consider $\mathrm{in}_{-e_1}(I_A)$. By Lemma 4.1, this initial ideal has an embedded prime. But since $\mathrm{in}_{-e_1}(I_A)$ is a monomial ideal, Theorem 2.4 implies that $\mathrm{in}_{-e_1}(I_A)$ has an embedded prime of the form $\langle \partial_j : j \notin \{1\} \cup \tau \rangle$, where $\tau \subset \{2, \dots, n\}$ has cardinality $d-2$. By interchanging the columns of A we may assume that $\tau = \{m+3, \dots, n\}$. Since standard pairs carry all the information about the associated primes of monomial ideals, we conclude that $\mathrm{in}_{-e_1}(I_A)$ has an embedded standard pair $(\partial^\nu, \{1\} \cup \tau)$.

Observation 4.3. We may assume that the set obtained from $P_\nu^{\overline{\{1\}\cup\tau}}$ by reversing the weight inequality $(B \cdot z)_1 \leq 0$ is bounded, or that the hyperplane $\{z : (B \cdot z)_r = 0\}$ coincides with the hyperplane $\{z : (B \cdot z)_1 = 0\}$ for some $r \notin \{1\} \cup \tau$. In the latter case, the first row of B is a negative multiple of the r-th row of B, and $\nu_r > 0$.

Proof. Suppose that the hyperplanes $\{z : (B \cdot z)_i = 0\}$, for $i \notin \{1\} \cup \tau$, are pairwise distinct. To check the assertion, first notice that, since $\langle \partial_j : j \notin \{1\} \cup \tau \rangle$ is an associated prime of $\mathrm{in}_{-e_1}(I_A)$, it is contained in a minimal prime $\langle \partial_j : j \notin \{1, l\} \cup \tau \rangle$ of this initial ideal. Here $l \in \{2, \dots, m+2\}$. Thus, the set obtained from $P_\nu^{\overline{\{1\}\cup\tau}}$ by removing the inequality $(B \cdot z)_l \leq \nu_l$ must be a simplex. Now, by Theorem 2.2, any lattice point $z \neq 0$ in this simplex satisfies $(B \cdot z)_l > \nu_l$, and such lattice points exist. Pick $z \in \mathbb{Z}^m \backslash \{0\}$ in our simplex such that $(B \cdot z)_l$ is minimal. Notice that the lattice point z is unique, since the weight vector $-e_l$ is generic. This follows from results in [14, Section 5].

Now let $\partial^\mu = \partial_1^{-(B\cdot z)_1 - 1} \prod_{i \in \{2, \dots, \hat{l}, \dots, m+2\}} \partial_i^{\nu_i - (B \cdot z)_i}$ and consider the pair $(\partial^\mu, \{l\} \cup \tau)$. It is easy to check that this pair satisfies the conditions of Theorem 2.2 for the weight vector $-e_l$. Thus, interchanging the first and l-th columns of A, and replacing ν by μ, we obtain a standard pair as we desired.

Finally, if the hyperplanes $\{z : (B \cdot z)_r = 0\}$ and $\{z : (B \cdot z)_s = 0\}$ coincide, for $r, s \notin \{1\} \cup \tau$, we use the above argument to change the weight vector $-e_1$ by the weight vector $-e_s$, and assume that $\{z : (B \cdot z)_r = 0\}$ and $\{z : (B \cdot z)_1 = 0\}$ are parallel. Of course, the first row of B is a multiple of the r-th row of B. If it were a positive multiple, then for each $v \in \ker_{\mathbb{Z}}(A)$, $\mathrm{in}_{-e_1}(\partial^{v_+} - \partial^{v_-})$ does not contain the variable ∂_r, which implies $r \in \tau$, a contradiction. Thus the first and r-th rows of B are negative multiples of each other. To see that $\nu_r > 0$, notice that $\nu_r = 0$ would contradict the last assertion of Theorem 2.2. $\square$

We are now ready to start the construction of our candidates for exceptional parameters. Pick an embedded standard pair $(\partial^\nu, \{1\} \cup \tau)$, and consider the set

$$\left\{ (\partial^\mu, \{1\} \cup \tau) : \begin{array}{c} (\partial^\mu, \{1\} \cup \tau) \in S(\mathrm{in}_{-e_1}(I_A)), \\ \exists y_\mu \in \mathbb{C}^m \text{ such that } \mu_i = \nu_i - (B \cdot y_\mu)_i \\ \text{for } 2 \leq i \leq m+2, \text{ and } 0 > (B \cdot y_\mu)_1 \in \mathbb{Z} \end{array} \right\}.$$

If this set is nonempty, select a standard pair $(\partial^\eta, \{1\} \cup \tau)$ such that $(B \cdot y_\eta)_1$ is maximal. Otherwise, rename ν to η. This choice implies the following fact.

Observation 4.4. If $(\partial^\mu, \{1\}\cup\tau)$ is a standard pair of $\mathrm{in}_{-e_1}(I_A)$ such that there exists $y \in \mathbb{C}^m$ that satisfies

1. $\mu_i = \eta_i - (B\cdot y)_i$ for $2 \leq i \leq m+2$, and
2. $(B\cdot y)_1 \in \mathbb{Z}$,

then $(B\cdot y)_1 \leq 0$.

Choose generic numbers α_i for $i \in \tau$. Let $\alpha = \sum_{i\in\tau}\alpha_i e_i$. We look at the fake exponents of $H_A(A\cdot(\eta+\alpha))$ with respect to the weight vector $-e_1$. Since the α_i are generic, we may assume that if u is a fake exponent and $u_1 = 0$, then $u_i \in \mathbb{Z}$ implies $i \notin \tau$. In particular, the numbers α_i are non integers. We define two sets:

$$F := \left\{ \begin{array}{c} \text{fake exponents of } H_A(A\cdot(\eta+\alpha)) \text{ with respect to } -e_1 \\ \text{that have minimal negative support} \end{array} \right\}$$

and

$$K := \{u \in F : u_1 = 0\}.$$

Proposition 4.5. *The following condition holds:*

$$\text{(2)} \qquad \textit{If } u\in K,\ v\in F, \textit{ and } v-u\in\mathbb{Z}^m, \textit{ then } 1\notin \mathrm{nsupp}(v).$$

Proof. Suppose that Condition 2 does not hold. Then there are $u \in K$, $v \in F$ such that $v-u\in\mathbb{Z}^m$, and $1 \in \mathrm{nsupp}(v)$. Since both u and v have minimal negative support (and their only integer coordinates are indexed by $i \notin \tau$ by the choice of α) we conclude that $\mathrm{nsupp}(u) = \{s\}$ for some $2 \leq s \leq m+2$, and $\mathrm{nsupp}(v) = \{1\}$.

Suppose that the standard pair that gives rise to v is top-dimensional. Then it is of the form $(\partial^\mu, \{1,l\}\cup\tau)$ for some $2 \leq l \leq m+2$. Since $u_s < 0$, the standard pair corresponding to u is of the form $(\partial^\nu, \{1,s\}\cup\tau)$.

Suppose $l = s$. Remember that we have $z \in \mathbb{Z}^m$ such that $v - B\cdot z = u$. Then $(B\cdot z)_j \leq v_j = \mu_j$ for $j \notin \{1,l\}\cup\tau$, and $(B\cdot z)_1 < 0$. Then $z \in \overline{P_\mu^{\{1,l\}\cup\tau}} \cap \mathbb{Z}^m = \{0\}$, a contradiction. Thus $l \neq s$.

We consider two cases, as in the proof of Observation 4.3. If two of the hyperplanes $\{y : (B\cdot y)_i = 0\}$, $i \notin \tau$, coincide, then $\{y : (B\cdot y)_r = 0\}$ is equal to $\{y : (B\cdot y)_1 = 0\}$ for some $2 \leq r \leq m+2$. This implies that $l = s = r$, because otherwise the sets $\overline{P_\nu^{\{1,s\}\cup\tau}}$ and $\overline{P_\mu^{\{1,l\}\cup\tau}}$ are unbounded. This contradicts the previous paragraph.

Now suppose that the hyperplanes $\{y : (B\cdot y)_i = 0\}$, $i \notin \tau$, are pairwise distinct. Since $\overline{P_\nu^{\{1,s\}\cup\tau}}$ and $\overline{P_\mu^{\{1,l\}\cup\tau}}$ are simplices, reversing the s-th or the l-th inequality in our original polytope $\overline{P_\eta^{\{1\}\cup\tau}}$ yields bounded sets. By Observation 4.3, if we reverse the inequality $(B\cdot z)_1 \leq 0$ we also get a bounded set. In view of Lemma 2.5, we derive a contradiction.

In conclusion, the standard pair corresponding to v cannot be top-dimensional, so it must be of the form $(\partial^\mu, \{1\}\cup\tau)$, by our choice of α. However, by Observation 4.4, a fake exponent of $H_A(A\cdot(\eta+\alpha))$ coming from such a standard pair cannot have a negative integer first coordinate. This contradiction concludes the proof. $\square$

In what follows, we will prove the following version of Theorem 1.6.

Theorem 4.6. *Suppose* $\mathrm{in}_{-e_1}(I_A)$ *is a monomial ideal with an embedded standard pair* $(\partial^\eta, \{1\} \cup \tau)$ *as in Observations 4.3 and 4.4, where* $\tau = \{m+3, \ldots, n\}$*. Choose generic numbers* α_i *as above. Then*

$$\beta := A \cdot (\eta + \alpha - e_1) \in \mathcal{E}(A).$$

Remark 4.7. *(Theorem 1.6 follows from Theorem 4.6.)* Notice that, once Theorem 4.6 is proved, Proposition 3.8 will lift the assumption that the numbers α_i are generic. Thus we will have produced a $(d-2)$-dimensional affine space contained in $\mathcal{E}(A)$, and the proof of Theorem 1.6 will be complete.

Remark 4.8. From now on, we work under the hypotheses and notation of Theorem 4.6.

We want to show that $\mathrm{rank}\,(H_A(\beta)) > \mathrm{vol}\,(A)$. One way to do this is to show that $\mathrm{rank}\,(H_A(\beta)) > \mathrm{rank}\,(H_A(A \cdot (\eta + \alpha)))$, since $\mathrm{rank}\,(H_A(A \cdot (\eta + \alpha))) \geq \mathrm{vol}\,(A)$. The tool to compare these two ranks is the D-module map (see [13, Section 4.5])

$$D/H_A(\beta) \longrightarrow D/H_A(A \cdot (\eta + \alpha))$$

given by right multiplication by the operator ∂_1. This induces a vector space homomorphism between the solution spaces of $H_A(A \cdot (\eta+\alpha))$ and $H_A(\beta)$: if φ is a solution of $H_A(A \cdot (\eta+\alpha))$, then $\partial_1 \varphi$ (the derivative of φ with respect to the variable x_1) is a solution of $H_A(\beta)$. It is this vector space map that we want to study. More precisely, we want information about the dimension of its kernel and cokernel.

5. The kernel of the map ∂_1

We start our analysis of the map ∂_1 between the solution spaces of $H_A(A \cdot (\eta + \alpha))$ and $H_A(\beta)$ by describing its kernel. The following proposition is the first step in this direction.

Lemma 5.1. *If* $u \in K$*, then the canonical series corresponding to* u*,* ϕ_u *equals the monomial* x^u*. Consequently,* $\partial_1 \phi_u = 0$*. Conversely, if* ϕ_u *is a logarithm-free canonical series such that* $\partial_1 \phi_u = 0$*, then* $u \in K$*.*

Proof. To see that ϕ_u is a monomial, it is enough to show that $N_u = \{0\}$. Remember that $N_u = \{B \cdot z : z \in \mathbb{Z}^m \text{ and } \mathrm{nsupp}(u - B \cdot z) = \mathrm{nsupp}(u)\}$. Since u is a fake exponent with respect to $-e_1$, there is a standard pair $(\partial^\mu, \{1\} \cup \sigma)$ of $\mathrm{in}_{-e_1}(I_A)$ such that $u_i = \mu_i \in \mathbb{N}$ for $i \notin \{1\} \cup \sigma$. Pick $z \in \mathbb{Z}^m$ such that $B \cdot z \in N_u$. Then, since $\mathrm{nsupp}(u - B \cdot z) = \mathrm{nsupp}(u)$, $(B \cdot z)_i \leq u_i = \mu_i$ for $i \notin \{1\} \cup \sigma$, and $(B \cdot z)_1 \leq u_1 = 0$. This means that $z \in P_\mu^{\overline{\{1\} \cup \sigma}} \cap \mathbb{Z}^m$, so that, by Theorem 2.2, $z = 0$. The rest of the assertions are trivial. $\square$

Theorem 5.2. $\ker\,(\partial_1) = \mathrm{Span}\,\big\{x^u : u \in K\big\}$.

Proof. Let φ be a (possibly logarithmic) solution of $H_A(A \cdot (\eta + \alpha))$ such that $\partial_1 \varphi = 0$. The function φ is a linear combination of canonical series with respect to $-e_1$. Write $\varphi = \varphi_1 + \cdots + \varphi_l$ where each φ_i is a linear combination of canonical series whose exponents differ by integer vectors, and the exponents in φ_i and φ_j do not differ by integers if $i \neq j$.

It is clear that $\partial_1 \varphi_i = 0$ for $1 \leq i \leq l$, so we can reduce to the case when φ is a linear combination of canonical series solutions whose exponents differ by integer vectors, and we assume this from now on.

Write φ in the form of Observation 3.6 for some $\delta \in \mathcal{S}_{\max}$, so that

$$\varphi = \varphi_\delta + f_\delta \log(x)^\delta$$

where φ_δ contains only logarithmic terms that less than or incomparable to δ. Since $\partial_1 \varphi = 0$, f_δ is a logarithm-free solution of $H_A(A \cdot (\eta + \alpha))$ that is constant with respect to x_1. Thus, f_δ is a linear combination of logarithm-free canonical series whose corresponding fake exponents differ by integer vectors, and have first coordinate equal to zero. This means that f_δ is a linear combination of functions $\phi_{u^{(i)}} = x^{u^{(i)}}$ with $u^{(i)} \in K$, differing pairwise by integer vectors.

Now rewrite the function φ in the form of Proposition 3.7, that is $\varphi = \sum x^v g_v(\log(x))$. By the previous reasoning, we can choose $u \in K$ such that $g_u \neq 0$. We will show that g_u is a constant.

Let $(\partial^\mu, \{1\} \cup \sigma)$ be a standard pair of $\operatorname{in}_{-e_1}(I_A)$ such that $u_i = \mu_i$ for $i \notin \{1\} \cup \sigma$. By our choice of α, $u_i \in \mathbb{Z}$ implies $i \notin \tau$, so that $\sigma \supseteq \tau = \{m+3, \dots, n\}$.

We first consider the case when $\sigma = \tau$. Pick $i \notin \tau$. By Theorem 2.2 applied to $(\partial^\mu, \{1\} \cup \tau)$, there exists $z \in \mathbb{Z}^m$ such that $(B \cdot z)_j \leq \mu_j$ for $j \notin \{1, i\} \cup \tau$, $(B \cdot z)_i > \mu_i$ and $(B \cdot z)_1 < 0$. Now choose a term t^γ appearing with nonzero coefficient in $g_u(t_1, \dots, t_n)$ such that t_i appears to a maximal power (among the terms in g_u). Since $\partial_1 \varphi = 0$,

$$0 = \partial^{(B \cdot z)_-} \varphi = \partial^{(B \cdot z)_+} \varphi.$$

Now, φ contains a summand that is a nonzero multiple of $x^u \log(x)^\gamma$. If $\gamma_i \neq 0$, then $\partial^{(B \cdot z)_+} \varphi$ contains a term that is a nonzero multiple of

$$\frac{\left(\partial^{(B \cdot z)_+ - ((B \cdot z)_i - u_i) e_i} x^u\right) \log(x)^{\gamma - e_i}}{x_i^{(B \cdot z)_i - u_i}} \quad . \tag{3}$$

By the construction of z, this term is nonzero. However, it cannot be cancelled with any other term from $\partial^{(B \cdot z)_+} \varphi = 0$, because such a term would have to come from g_u, and γ_i was chosen maximal among those terms. This contradiction implies that $\gamma_i = 0$. We conclude that g_u is constant with respect to the i-th variable, and this is true for all $i \notin \{1\} \cup \tau$.

However, by Proposition 3.7

$$g_u(t_1, \dots, t_n) = c_0 + c^{(1,1)} \cdot (t_1, \dots, t_n) + \cdots + \prod_{k=1}^{l} c^{(l,k)} \cdot (t_1, \dots, t_n),$$

where $c_0 \in \mathbb{C}$, the vectors $c^{(j,k)}$ belong to the kernel of A, and

$$c^{(j,k)} \cdot (t_1, \ldots, t_n) = c_1^{(j,k)} t_1 + \cdots + c_n^{(j,k)} t_n \ .$$

Since g_u is constant with respect to the i-th variable for all $i \notin \{1\} \cup \tau$, $c_i^{(j,k)} = 0$ for $i \notin \{1\}\tau$, $j = 1, \ldots l$, $1 \leq k \leq j$. By Lemma 2.3, m of the rows of B indexed by $i \notin \{1\} \cup \tau$ are linearly independent. This implies that all the vectors $c^{(j,k)}$ must be zero, and we conclude that g_u is constant.

Now we need to show that g_u is constant in the case when the standard pair corresponding to u is $(\partial^\mu, \{1\} \cup \sigma)$, and σ strictly contains τ. In this case, $\sigma = \tau \cup \{r\}$, for some $1 < r \leq m+2$. Clearly, if $u_r \notin \mathbb{N}$, the previous arguments still apply, because no matter what $(B \cdot z)_r$ is, the term (3) cannot vanish. Thus, we may assume that $u_r \in \mathbb{N}$.

The polytope $P_\mu^{\overline{\{1\} \cup \sigma}}$ is a simplex, and $0 \in P := P_\mu^{\overline{\{1\} \cup \sigma}} \cap \{z \in \mathbb{R}^m : (B \cdot z)_r \leq u_r\}$. This means that Lemma 2.5 applies, and so there is a set $T \subset \{1, \ldots, \hat{r}, \ldots, m+2\}$ of cardinality m such that, for each $i \in T$, the set obtained from P by reversing the inequality $(B \cdot z)_i \leq u_i$ is unbounded. Notice that, by Observation 4.3, we may assume that $1 \notin T$.

If $i \in T$, we can choose $z \in \mathbb{Z}^m$ such that $(B \cdot z)_j \leq u_j$ for $j \notin \{1, i\} \cup \tau$, $(B \cdot z)_i > u_i$, and $(B \cdot z)_1 < 0$. Now we can repeat what we did before, and conclude that g_u is constant with respect to the i-th variable, for all $i \in T$. Since the rows of B indexed by T are linearly independent, using Proposition 3.7, we conclude that g_u is constant.

Now, remember that $g_u(\log(x))$ contains a term $\log(x)^\delta$ that is maximal among the logarithmic terms appearing in φ. Since g_u is constant, $\delta = 0$, which implies that φ is logarithm-free. But then it is clear that φ belongs to Span $\{x^u : u \in K\}$. ◻

6. The cokernel of the map ∂_1

In this section we produce a subspace of coker (∂_1) whose dimension equals the dimension of ker (∂_1).

Lemma 6.1. *For each $u \in K$, the vector $u - e_1$ is a fake exponent with minimal negative support of $H_A(\beta)$.*

Proof. Pick $u \in K$ corresponding to a standard pair $(\partial^\mu, \{1\} \cup \sigma)$. Clearly, $u - e_1$ is the fake exponent of $H_A(\beta)$ corresponding to this standard pair. Suppose that $u - e_1$ does not have minimal negative support. Then we can find $z \in \mathbb{Z}^m$ such that $u - e_1 - B \cdot z$ has minimal negative support strictly contained in nsupp$(u - e_1)$. In particular, $z \neq 0$. Notice that $1 \in$ nsupp$(u - e_1 - B \cdot z)$, because otherwise $z \in P_\mu^{\overline{\{1\} \cup \sigma}} \cap \mathbb{Z}^m = \{0\}$. This implies that the functional $-e_1$ is minimized in the set $\{B \cdot y : y \in \mathbb{Z}^m,$ and nsupp$(u - e_1 - B \cdot z - B \cdot y) =$ nsupp$(u - e_1 - B \cdot z)\}$. This minimum must be unique since $-e_1$ is a generic weight vector. As a consequence, we may assume that $u - e_1 - B \cdot z$ is a fake exponent of $H_A(\beta)$ with respect to the weight vector $-e_1$. Then $u - B \cdot z$ is a fake exponent of $H_A(A \cdot (\eta + \alpha))$, it has minimal negative support, and $(u - B \cdot z)_1 < 0$. The last two facts follow since u has minimal negative support. But now this contradicts Proposition 4.5. ◻

Theorem 6.2. *The set* $\operatorname{Span}\{\phi_{u-e_1} : u \in K\}$ *does not intersect the image of the map* ∂_1*.*

Proof. Suppose there is a solution ψ of $H_A(A\cdot v)$ such that $\partial_1\psi$ lies in $\operatorname{Span}\{\phi_{u-e_1} : u \in K\}$. As in the proof of Theorem 5.2, we may assume that ψ is a linear combination of canonical series whose exponents differ by integer vectors.

We proceed as in the part of the proof of Theorem 5.2 where we show that the functions φ_i are logarithm-free. The first step is to write $\psi = \psi_\delta + f_\delta \log(x)^\delta$ for every $\delta \in \mathcal{S}_{\max}$ as in Observation 3.6. The function f_δ belongs to the kernel of ∂_1, so we may choose $u \in K$ such that x^u appears with a nonzero coefficient in f_δ. Now we rewrite $\psi = \sum_v x^v g_v(\log(x))$ as in Proposition 3.7. We want to compute the polynomial g_u.

As in Theorem 5.2 we consider two cases, according to whether the standard pair $(\partial^\mu, \{1\}\cup\sigma)$ corresponding to the fake exponent u is embedded or not.

In the first case, $\sigma = \tau$, and $\operatorname{nsupp}(u) = \emptyset$. Choose $i \notin \{1\}\cup\tau$ and a vector $z \in \mathbb{Z}^m$ such that $(B\cdot z)_j \leq \mu_j$ for $j \notin \{1,i\}\cup\tau$, $(B\cdot z)_i > \mu_i$, and $(B\cdot z)_1 < 0$. We can do this by Theorem 2.2.

Now choose a term t^γ appearing with nonzero coefficient in the polynomial $g_u(t_1,\dots,t_n)$ such that t_i appears to a maximal power (among the terms in g_u), and consider the term (3) from the proof of Theorem 5.2:

$$\frac{\left(\partial^{(B\cdot z)_+ - ((B\cdot z)_i - u_i)e_i} x^u\right)\log(x)^{\gamma - e_i}}{x_i^{(B\cdot z)_i - u_i}}\ ,$$

which appears with a nonzero coefficient in $\partial^{(B\cdot z)_+}\psi = \partial^{(B\cdot z)_-}\psi$. Since $(B\cdot z)_1 < 0$, this function is a further derivative of $\partial_1\psi \in \operatorname{Span}\{\phi_{u-e_1} : u \in K\}$.

As $\operatorname{nsupp}(u) = \emptyset$, there are no fake exponents $v \in K$ such that $u - v \in \mathbb{Z}^m$. Thus $\partial_1\psi$ must be a multiple of ϕ_{u-e_1}, which has neither logarithmic terms, nor terms which contain a strictly negative integer power of x_i. We conclude that $\gamma_i = 0$. The same argument as in Theorem 5.2 now implies that g_u is constant. This is a contradiction because x^{u-e_1} can only appear in $\partial_1\psi$ if ψ has a term $x^u\log(x_1)$.

Now we need to consider the case when the standard pair $(\partial^\mu, \{1\}\cup\sigma)$ is top-dimensional, that is $\sigma = \tau\cup\{r\}$, for some $2 \leq r \leq m+2$. If $u_r \notin \mathbb{Z}$, the same argument as above shows that g_u is constant. If $u_r \in \mathbb{N}$, we need to combine the previous argument with the reasoning from Theorem 5.2 to again conclude that g_u is constant. Finally if u_r is a negative integer, we observe that $\partial_1\psi$ is a linear combination of functions ϕ_{v-e_1}, where the vectors v belong to K and differ with u by an integer vector. But then, for each $i \notin \{1,r\}\cup\tau$, a further derivative of $\partial_1\psi$ has no logarithmic terms, or terms that contain strictly negative integer powers of the variable x_i and x_r. Unless g_u is constant with respect to the i-th variable, this contradicts the fact that $\partial^{(B\cdot z)_+}\psi$ contains a nonzero multiple of the term (3). As before, we conclude that g_u is constant, a contradiction. ∎

7. Producing a rank jump

We now analyze what solutions of $H_A(A\cdot(\eta+\alpha))$ and $H_A(\beta)$ arise from the standard pair $(\partial^\eta, \{1\}\cup\tau)$. Clearly, $\eta+\alpha$ is the fake exponent of $H_A(A\cdot(\eta+\alpha))$ corresponding to this standard pair. By Lemma 5.1, the corresponding A-hypergeometric series is $\phi_{\eta+\alpha} = x^{\eta+\alpha}$.

For each $2 \leq i \leq m+2$, that is, for each $i \notin \{1\}\cup\tau$, we can choose $z^{(i)} \in \mathbb{Z}^m$ such that $(B\cdot z^{(i)})_j \leq \eta_j$ for $j \notin \{1,i\}\cup\tau$, $(B\cdot z^{(i)})_i > \eta_i$ and $(B\cdot z^{(i)})_1 < 0$. Let $z^{(1)} = 0$, and consider the vectors $\eta+\alpha-e_1-B\cdot z^{(i)}$. We have $\operatorname{nsupp}(\eta+\alpha-e_1-B\cdot z^{(i)}) = \{i\}$. By Lemma 6.1, $\eta+\alpha-e_1$ has minimal negative support. It follows that all the $\eta+\alpha-e_1-B\cdot z^{(i)}$ have minimal negative support. Of course, this is not enough to guarantee that they all give rise to logarithm-free solutions of $H_A(\beta)$. However, some of them do, for instance $\eta+\alpha-e_1$, since it is a fake exponent. The corresponding A-hypergeometric series is $\phi_1 := \phi_{\eta+\alpha-e_1}$. The purpose of the next proposition is to show that there is at least one other solution of $H_A(\beta)$ arising this way.

Proposition 7.1. *There exists at least one $i \notin \{1\}\cup\tau$ for which we can find a vector $y^{(i)} \in \mathbb{Z}^m$ such that $\eta+\alpha-e_1-B\cdot(z^{(i)}+y^{(i)})$ is a fake exponent with minimal negative support equal to $\{i\}$.*

Proof. Let $N_i = \{B\cdot z : z \in \mathbb{Z}^m \text{ and } \operatorname{nsupp}(u-e_1-B\cdot z) = \{i\}\}$ for $i \notin \{1\}\cup\tau$. If we can show that the linear functional $-e_1$ is minimized uniquely in the set $U_i = \{\eta+\alpha-B\cdot z : B\cdot z \in N_i\}$, then this set will contain a fake exponent, which is what we want to prove. Since the weight vector $-e_1$ is generic, it is enough to see that $-e_1$ is minimized in U_i, or equivalently, that it is maximized in N_i.

As in the proof of Observation 4.3 there are two cases. Either the hyperplanes $\{z : (B\cdot z)_i = 0\}$ are all distinct, for $i \notin \tau$, or two of those coincide.

In the second case, we know that $\{z : (B\cdot z)_1 = 0\}$ equals $\{z : (B\cdot z)_r = 0\}$ for a certain $2 \leq r \leq m+2$. But then, for each $i \notin \{1,r\}\cup\tau$, $B\cdot z \in N_i$ implies $(B\cdot z)_r \leq \eta_r$. Since the r-th row of B is a negative multiple of the first row, the linear functional $-e_1$ is bounded above in N_i. Thus it is maximized in N_i, and we obtain logarithm-free canonical solutions ϕ_i of $H_A(\beta)$ corresponding to $\eta+\alpha-e_1-B\cdot z^{(i)}$ for $i \notin \{1,r\}\cup\tau$.

Now assume that the hyperplanes $\{z : (B\cdot z)_i = 0\}$ are pairwise different for $i \notin \tau$, and remember that $\langle \partial_i : i \notin \{1\}\cup\tau\rangle$ is an embedded prime of $\operatorname{in}_{-e_1}(I_A)$. Then there is a minimal prime of this initial ideal of the form $\langle \partial_i : i \notin \{1,l\}\cup\tau\rangle$, for some $l \notin \{1\}\cup\tau$. Since there must be standard pairs corresponding to this minimal prime, we conclude that the set obtained from $P_\eta^{\overline{\{1\}\cup\tau}}$ by deleting the inequality $(B\cdot z)_l \leq \eta_l$ is a simplex. Thus N_l is bounded, so that $-e_1$ will be maximized in N_l, and we obtain ϕ_l, the corresponding solution of $H_A(\beta)$. Notice that, by Lemma 2.5, the linear functional $-e_1$ attains no maximum in N_i for $i \neq 1, l$. $\square$

The following theorem is the fundamental ingredient to produce a rank jump.

Theorem 7.2. *Let $s \geq 2$ be the dimension of the span of the functions constructed in Proposition 7.1 (this includes the function $\phi_1 = \phi_{\eta+\alpha-e_1}$). The intersection of this span with the image of ∂_1 has dimension at most $s-2$.*

Proof. First notice that if ψ is a logarithm-free solution of $H_A(A\cdot(\eta+\alpha))$, and $\partial_1\psi$ lies in the span of the functions from Proposition 7.1, then ψ is a linear combination of canonical, logarithm-free solutions of $H_A(A\cdot(\eta+\alpha))$ that are either constant with respect to x_1, or whose fake exponents differ by an integer vector with $\eta+\alpha$. Since $\mathrm{nsupp}(\eta+\alpha)=\emptyset$, the only such fake exponent is $\eta+\alpha$. Thus $\partial_1\psi = 0$.

Now let ψ be a logarithmic solution of $H_A(A\cdot(\eta+\alpha))$ such that $\partial_1\psi$ lies in the span of the functions constructed in Proposition 7.1. We may assume that ψ does not contain a summand x^u, $u \in K$. As in the proof of Theorem 6.2, we can write $\psi = \sum x^v g_v(\log(x))$, where $g_{\eta+\alpha}(\log(x))$ contains all the maximal logarithmic terms.

Pick $i \notin \{1\}\cup\tau$, and remember the vector $z^{(i)} \in \mathbb{Z}^m$ from the paragraph before Proposition 7.1. An argument similar to that in the proof of Theorem 6.2 shows that, if $g_{\eta+\alpha}(\log(x))$ contains a power of $\log(x_i)$ greater than 1, then $\partial_1\psi$ cannot be logarithm-free. This and Proposition 3.7 imply that $g_{\eta+\alpha}(\log(x)) = c_1\log(x_1)+\cdots+c_n\log(x_n)$, where the vector $(c_1,\dots,c_n)$ lies in the kernel of A. Thus,

$$\psi = \psi_{\eta+\alpha} + x^{\eta+\alpha}(c_1\log(x_1)+\cdots+c_n\log(x_n)) \ ,$$

where $\psi_{\eta+\alpha}$ is logarithm-free, and contains no terms x^u, $u\in K$. Notice that, if $\eta+\alpha-e_1+(B\cdot z)_i$ does not give rise to a solution of $H_A(\beta)$ as in Proposition 7.1, then $c_i=0$. This follows from the same arguments that proved Theorem 6.2.

We claim that, once the numbers $c_1,\dots,c_n$ are fixed, the function ψ itself is fixed. This is because the difference between two such functions would be a logarithm-free solution of $H_A(A\cdot(\eta+\alpha))$ whose first derivative lies in the span of the functions from Proposition 7.1. By the first paragraph of this proof, this implies that our difference lies in the kernel of ∂_1. But since it cannot contain terms in x^u, $u\in K$, we conclude that this difference must be zero.

Now we consider two cases, as in the proof of Proposition 7.1. In the case that a hyperplane $\{z : (B\cdot z)_r = 0\}$ coincides with $\{z : (B\cdot z)_1 = 0\}$, we have $m+1$ linearly independent functions ϕ_i, $i\notin\{r\}\cup\tau$ from Proposition 7.1. But if $\partial_1\psi$ lies in the span of these functions, we know that

$$\psi = \psi_{\eta+\alpha} + x^{\eta+\alpha}(c_1\log(x_1)+\cdots+c_n\log(x_n)) \ ,$$

where $(c_1,\dots,c_n)$ belongs to the kernel of A, and $c_r=0$. Since the dimension of the kernel of A is m, we conclude that the dimension of the space of such functions ψ is at most $m-1$. Thus, the intersection of the image of ∂_1 and the span of the functions ϕ_i, $i\notin\{r\}\cup\tau$ has dimension at most $m-1=(m+1)-2=s-2$.

In the case when all the hyperplanes $\{z : (B\cdot z)_i = 0\}$, $i\notin\tau$ are distinct, we have only two functions, ϕ_1 and ϕ_l from Proposition 7.1. Thus, if $\partial_1\psi$ lies in their span,

$$\psi = \psi_{\eta+\alpha} + x^{\eta+\alpha}(c_1\log(x_1)+\cdots+c_n\log(x_n)) \ ,$$

where $(c_1, \dots, c_n)$ belongs to the kernel of A, and $c_i = 0$, for $i \notin \{1, l\} \cup \tau$. By Lemma 2.3, this implies that $(c_1, \dots, c_n)$ is the zero vector, so that ψ vanishes. Thus the intersection of $\mathrm{Span}\,\{\phi_1, \phi_l\}$ with the image of ∂_1 is $\{0\}$, which has dimension $s-2$. ☐

We are finally ready to prove Theorem 4.6.

Proof of Theorem 4.6. In Theorem 6.2 we produced one function in the cokernel of ∂_1 for each function in the kernel of ∂_1. Furthermore, in Theorem 7.2, we produced at least 2 linearly independent functions in the cokernel of ∂_1 corresponding to $x^{\eta+\alpha}$. All these functions are clearly linearly independent in the cokernel of ∂_1. This means that:

$$\dim(\mathrm{coker}\,(\partial_1)) \geq \dim(\ker\,(\partial_1)) + 1\ .$$

Adding $\dim(\mathrm{im}\,(\partial_1))$ to both sides of this inequality we obtain:

$$\dim(\mathrm{coker}\,(\partial_1)) + \dim(\mathrm{im}\,(\partial_1)) \geq \dim(\ker\,(\partial_1)) + \dim(\mathrm{im}\,(\partial_1)) + 1\ ,$$

or equivalently,

$$\mathrm{rank}\,(H_A(\beta)) \geq \mathrm{rank}\,(H_A(A \cdot v)) + 1\ .$$

Since $\mathrm{rank}\,(H_A(A \cdot v)) \geq \mathrm{vol}\,(A)$, it follows that

$$\mathrm{rank}\,(H_A(\beta)) \geq \mathrm{vol}\,(A) + 1\ .$$

☐

8. An example

Even though generic toric ideals are common among all toric ideals, it is hard in practice to construct examples of generic configurations A such that $\mathrm{vol}\,(A)$ is small. This is a disadvantage when we want to perform rank computations using computer algebra systems. In this section, we work out an example where I_A is *not* generic, although the techniques that proved Theorem 4.6 are still successful. Let

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 2 & 1 & 2 & 3 & 0 \\ 0 & 2 & 2 & 0 & 1 & 1 \end{pmatrix}\ .$$

The toric ideal is:

$$I_A = \left\langle \begin{array}{c} \partial_3\partial_4 - \partial_5\partial_6, \partial_1\partial_2 - \partial_5\partial_6, \partial_3^3\partial_5 - \partial_2^3\partial_6, \partial_1\partial_3^2\partial_5 - \partial_2^2\partial_4\partial_6 \\ \partial_1^2\partial_3\partial_5 - \partial_2\partial_4^2\partial_6, \partial_1^3\partial_5 - \partial_4^3\partial_6, \partial_2\partial_4^3 - \partial_1^2\partial_5^2, \partial_2^2\partial_4^2 - \partial_1\partial_3\partial_5^2 \\ \partial_2^3\partial_4 - \partial_3^2\partial_5^2, \partial_1\partial_3^3 - \partial_2^2\partial_6^2, \partial_1^2\partial_3^2 - \partial_2\partial_4\partial_6^2, \partial_1^3\partial_3 - \partial_4^2\partial_6^2 \end{array} \right\rangle\ .$$

Notice that $\mathrm{vol}\,(A) = 8$, and that I_A is not Cohen Macaulay. We compute the initial ideal of I_A with respect to $-e_1$:

$$\mathrm{in}_{-e_1}(I_A) = \left\langle \begin{array}{c} \partial_5\partial_6, \partial_3\partial_4 - \partial_5\partial_6, \partial_4^2\partial_6^2, \partial_2\partial_4\partial_6^2, \partial_2^2\partial_6^2, \partial_4^3\partial_6, \partial_2\partial_4^2\partial_6 \\ \partial_2^2\partial_4\partial_6, \partial_3^3\partial_5 - \partial_2^3\partial_6, \partial_2\partial_4^3, \partial_2^2\partial_4^2, \partial_2^3\partial_4 - \partial_3^2\partial_5^2 \end{array} \right\rangle\ .$$

Since this ideal is not a monomial ideal, I_A is not generic. However, we can try to look at its embedded primes, in the spirit of the generic case. Using the computer algebra system

`Singular`, we can compute the set of associated primes of $\mathrm{in}_{-e_1}(I_A)$. We obtain the following set:

$$\left\{ \begin{array}{c} \langle\partial_6,\partial_5,\partial_4\rangle, \langle\partial_6,\partial_4,\partial_3\rangle, \langle\partial_6,\partial_3,\partial_2\rangle, \langle\partial_5,\partial_4,\partial_2\rangle \\ \langle\partial_6,\partial_4,\partial_3,\partial_2\rangle, \langle\partial_5,\partial_4,\partial_3,\partial_2\rangle, \langle\partial_2,\partial_3,\partial_4,\partial_5,\partial_6\rangle \end{array} \right\}.$$

Corresponding to the ideal $\langle\partial_5,\partial_4,\partial_3,\partial_2\rangle$, $\mathrm{in}_{-e_1}(I_A)$ has a standard pair $(\partial_4,\{1,6\})$, according to Altmann's generalization of this notion (see [2]). We can try to use the construction from Theorem 4.6 applied to this standard pair. We obtain a line

$$\{A\cdot(-1,0,0,1,0,\alpha) = (\alpha,1,\alpha) : \alpha \in \mathbb{C}\}$$

of candidates for exceptional parameters.

We can check, using the computer algebra system `Macaulay2`, that $\operatorname{rank}(H_A(0,1,0)) = 9$. The calculation, performed in a `Linux` computer with dual Pentium III-700/100 processors and 512 MB of RAM, lasted about 30 minutes.

We would like to use the methods from the previous section to prove that $(\alpha,1,\alpha)$ is an exceptional parameter. The first difficulty we encounter is that we cannot use the weight vector $-e_1$ to compute canonical series, since this weight vector is not generic. To bypass this disadvantage, we consider the following initial ideal of $\mathrm{in}_{-e_1}(I_A)$:

$$\mathrm{in}_w(I_A) = \left\langle \begin{array}{c} \partial_5\partial_6, \partial_4\partial_3, \partial_4^2\partial_6^2, \partial_2\partial_4\partial_6^2, \partial_2^2\partial_6^2, \partial_4^3\partial_6 \\ \partial_2\partial_4^2\partial_6, \partial_2^2\partial_4\partial_6, \partial_5\partial_3^3, \partial_2\partial_4^3, \partial_2^2\partial_4^2, \partial_2^3\partial_4 \end{array} \right\rangle.$$

The set of standard pairs of this initial ideal is:

$$\begin{array}{ll} (\partial_3,\{1,2,3\}), & (1,\{1,4,5\}), \\ (1,\{1,2,3\}), & (\partial_4,\{1,6\}), \\ (\partial_2,\{1,3,6\}), & (\partial_2\partial_4^2,\{1,5\}), \\ (1,\{1,3,6\}), & (\partial_2^2\partial_4,\{1,5\}), \\ (\partial_3^2,\{1,2,5\}), & (\partial_2\partial_4,\{1,5\}), \\ (\partial_3,\{1,2,5\}), & (\partial_4^2\partial_6,\{1\}), \\ (1,\{1,2,5\}), & (\partial_2\partial_4\partial_6,\{1\}). \end{array}$$

Thus, the standard pair $(\partial_4,\{1,6\})$, which we used to construct our candidate for exceptional parameter, is also a standard pair on $\mathrm{in}_w(I_A)$. We will use this (generic) weight vector to compute canonical series solutions of $H_A(A\cdot(0,0,0,1,0,\alpha)) = H_A((1+\alpha,2,\alpha))$ and $H_A(A\cdot(-1,0,0,1,0,\alpha)) = H_A((\alpha,1,\alpha))$. As in the proof of Theorem 4.6, we will use the map ∂_1 between the solution spaces of $H_A((1+\alpha,2,\alpha))$ and $H_A((\alpha,1,\alpha))$.

If we chose $\alpha \neq 1,4,2/3$, so that the only embedded standard pair that produces a fake exponent is $(\partial_4,\{1,6\})$, the fake exponents of $H_A((1+\alpha,2,\alpha))$ with respect to w, ordered

by their corresponding standard pairs, are:

$$\begin{array}{ll}
(\alpha/2+1/2, 2-\alpha, 3\alpha/2-5/2, 0, 0, 1), & (3\alpha, 0, 0, 1-3\alpha, \alpha, 0), \\
(\alpha/2+1, 1-\alpha, 3\alpha/2-1, 0, 0, 0), & (0, 0, 0, 1, 0, \alpha). \\
(3/2, 0, 1/2, 0, 0, \alpha-1), & - \\
(1, 1, -1, 0, 0, \alpha), & - \\
(\alpha+2/3, \alpha-1/3, 0, 0, 2/3-\alpha, 0), & - \\
(\alpha+1/3, \alpha-5/3, 1, 0, 4/3-\alpha, 0), & - \\
(\alpha, \alpha-3, 2, 0, 2-\alpha, 0). & -
\end{array}$$

Now we choose $\alpha \neq -1, -2, -2/3, -1/3, 0$, so that none of the fake exponents of $H_A((1+\alpha, 2, \alpha))$ have zero first coordinate except $(0,0,0,1,0,\alpha)$. This implies that $\ker(\partial_1) = \mathrm{Span}\,\{\phi = x_4 x_6^{\alpha}\}$.

In order to prove that $\mathrm{rank}\,(H_A((\alpha, 1, \alpha))) > \mathrm{vol}\,(A) = 8$, we need to produce at least two functions in the cokernel of ∂_1. The first step is to make α generic so that the only fake exponent of $H_A((1+\alpha, 2, \alpha))$ that differs with $(0,0,0,1,0,\alpha)$ by an integer vector is $(1,1,-1,0,0,\alpha)$. Now, the vectors $(-1,0,0,1,0,\alpha)$ and $(0,1,-1,0,0,\alpha)$ are fake exponents of $H_A((\alpha,1,\alpha))$ with minimal negative support. Arguments similar to those in Section 7 show that these two functions are linearly independent elements of $\mathrm{coker}\,(\partial_1)$. In conclusion,

$$\mathrm{rank}\,(H_A((\alpha,1,\alpha))) > \mathrm{rank}\,(H_A((1+\alpha,2,\alpha))) \geq \mathrm{vol}\,(A)\ .$$

Details on how to generalize this example can be found in [8].

9. The geometry of the exceptional set

The first basic question about the exceptional set is to determine exactly when it is nonempty. There are other open problems in this area. One of them is to determine whether or not this set is Zariski closed. It is not hard to show, however, that it is Zariski constructible, if we make use of comprehensive Gröbner bases.

Proposition 9.1. *The exceptional set of a homogeneous matrix A is Zariski constructible, that is, it can be expressed as a finite Boolean combination of Zariski closed sets.*

Proof. We will prove our claim by presenting an algorithm to compute $\mathcal{E}(A)$ for a given homogeneous matrix A. This algorithm will rely on Gröbner basis computations in the Weyl algebra.

In this proof we will think of β as a parameter vector, that is, a vector of indeterminates, instead of an element of $\mathbb{C}^d$. Thus, the A-hypergeometric system $H_A(\beta)$ will no longer be an ideal in the Weyl algebra D, but an ideal in the *parametric* Weyl algebra $D_{\mathbb{C}[\beta]}$, where the parameters β_i commute with the variables x_i and ∂_i. Given a vector in $\mathbb{C}^d$, the *specialization map* corresponding to this vector is the map from $D_{\mathbb{C}[\beta]}$ to D that replaces the parameters β_i by the coordinates of our vector. This allows us to introduce the concept of a comprehensive Gröbner basis of a given left ideal I in the parametric Weyl algebra. Informally, a set G is a *comprehensive left Gröbner basis* of I if, for every vector in $\mathbb{C}^d$, the specialization of G with respect to that vector is a Gröbner basis of the corresponding specialization of I.

We are now ready to describe an algorithm to compute $\mathcal{E}(A)$.

Input: A homogeneous matrix A.
Output: The exceptional set $\mathcal{E}(A)$.

1. Compute a comprehensive left Gröbner basis of $H_A(\beta)$ with respect to the weight vector whose first n entries are zeros, and whose last n entries are ones. Call this basis G.
2. There are only finitely many initial ideals $\text{in}_{(0,\ldots,0,1,\ldots,1)}(H_A(\beta))$ under specialization of the parameter β. These are ideals in the polynomial ring $\mathbb{C}[x_1,\ldots,x_n,s_1,\ldots,s_n]$, call them $I_1,\ldots,I_l$.
3. For each $1 \leq j \leq l$, the comprehensive Gröbner basis G yields finitely many polynomial conditions of the form $g(\beta) = 0$ or $g(\beta) \neq 0$ which characterize the (constructible) subset T_j of $\mathbb{C}^d$ such that if we specialize β to an element of T_j, the initial ideal of the specialized A-hypergeometric system is I_j.
4. For $1 \leq j \leq l$ compute
$$r_j = \dim_{\mathbb{C}(x_1,\ldots,x_n)} \left(\frac{\mathbb{C}(x_1,\ldots,x_n)[s_1,\ldots,s_n]}{\mathbb{C}(x_1,\ldots,x_n)[s_1,\ldots,s_n] \cdot I_j} \right).$$
5. $\mathcal{E}(A) = \cup_{j:r_j > \text{vol}\,(A)} T_j$.

By [13, Formula 1.26], the number r_j produced in the fourth step of the previous algorithm equals the rank of any D-ideal whose initial ideal with respect to $(0,\ldots,0,1,\ldots,1)$ is I_j. This justifies the description of $\mathcal{E}(A)$ from the last step of the algorithm. □

Comprehensive Gröbner bases were introduced by Volker Weispfenning in [15]. This article deals only with the commutative case, and contains an explicit algorithm for computing comprehensive Gröbner bases. We refer to [6] for a proof of the existence of these objects in the non commutative case. Here, the authors argue that comprehensive Gröbner bases can be constructed, but they do not provide an explicit algorithm. For this reason, it would be very desirable to have a method for computing the exceptional set that did not require comprehensive Gröbner bases.

Acknowledgments: I am very grateful to Bernd Sturmfels for many inspiring discussions, and unfailingly useful suggestions. I would also like to thank Diane Maclagan, for her comments on a previous version of this article.

References

[1] Alan Adolphson, *Hypergeometric functions and rings generated by monomials*, Duke Math. J. **73** (1994), no. 2, 269–290.

[2] Klaus Altmann, *The chain property for the associated primes of $\mathcal{A}$-graded ideals*, Math. Res. Lett. **7** (2000), no. 5-6, 565–575.

[3] Eduardo Cattani, Carlos D'Andrea, and Alicia Dickenstein, *The $\mathcal{A}$-hypergeometric system associated with a monomial curve*, Duke Math. J. **99** (1999), no. 2, 179–207.

[4] I. M. Gel′fand, A. V. Zelevinsky, and M. M. Kapranov, *Hypergeometric functions and toric varieties*, Funktsional. Anal. i Prilozhen. **23** (1989), no. 2, 12–26.

[5] Serkan Hoşten and Rekha R. Thomas, *The associated primes of initial ideals of lattice ideals*, Math. Res. Lett. **6** (1999), no. 1, 83–97.

[6] Heinz Kredel and Volker Weispfenning, *Parametric gröbner bases in rings of solvable type*, Proc. IV International Conference on Computer Algebra in Physical Research, Joint Institute for Nuclear Research, Dubna, USSR (World Scientific, Singapore), 1991, pp. 236–244.
[7] Laura Felicia Matusevich, *Rank jumps in codimension* 2 *A-hypergeometric systems*, To appear in the Journal of Symbolic Computation. See also `arXiv:math.CO/0009148`.
[8] ———, *Combinatorial aspects of hypergeometric functions*, Ph.D. thesis, University of California at Berkeley, 2002.
[9] Ezra Miller, Bernd Sturmfels, and Kohji Yanagawa, *Generic and cogeneric monomial ideals*, J. Symbolic Comput. **29** (2000), no. 4-5, 691–708, Symbolic computation in algebra, analysis, and geometry (Berkeley, CA, 1998).
[10] Irena Peeva and Bernd Sturmfels, *Generic lattice ideals*, J. Amer. Math. Soc. **11** (1998), no. 2, 363–373.
[11] ———, *Syzygies of codimension* 2 *lattice ideals*, Math. Z. **229** (1998), no. 1, 163–194.
[12] Mutsumi Saito, *Logarithm-free A-hypergeometric series*, `arXiv:math.AG/0012257`.
[13] Mutsumi Saito, Bernd Sturmfels, and Nobuki Takayama, *Gröbner deformations of hypergeometric differential equations*, Springer-Verlag, Berlin, 2000.
[14] Bernd Sturmfels, Robert Weismantel, and Günter M. Ziegler, *Gröbner bases of lattices, corner polyhedra, and integer programming*, Beiträge Algebra Geom. **36** (1995), no. 2, 281–298.
[15] Volker Weispfenning, *Comprehensive Gröbner bases*, J. Symbolic Comput. **14** (1992), no. 1, 1–29.
[16] Günter M. Ziegler, *Lectures on polytopes*, Springer-Verlag, New York, 1995.

University of California at Berkeley
E-mail address: `laura@math.berkeley.edu`